



# Mellanox ConnectX<sup>®</sup>-4 Lx Firmware Release Notes

---

Rev 14.26.1040

NOTE:

THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT ("PRODUCT(S)") AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES "AS-IS" WITH ALL FAULTS OF ANY KIND AND SOLELY FOR THE PURPOSE OF AIDING THE CUSTOMER IN TESTING APPLICATIONS THAT USE THE PRODUCTS IN DESIGNATED SOLUTIONS. THE CUSTOMER'S MANUFACTURING TEST ENVIRONMENT HAS NOT MET THE STANDARDS SET BY MELLANOX TECHNOLOGIES TO FULLY QUALIFY THE PRODUCT(S) AND/OR THE SYSTEM USING IT. THEREFORE, MELLANOX TECHNOLOGIES CANNOT AND DOES NOT GUARANTEE OR WARRANT THAT THE PRODUCTS WILL OPERATE WITH THE HIGHEST QUALITY. ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT ARE DISCLAIMED. IN NO EVENT SHALL MELLANOX BE LIABLE TO CUSTOMER OR ANY THIRD PARTIES FOR ANY DIRECT, INDIRECT, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING, BUT NOT LIMITED TO, PAYMENT FOR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY FROM THE USE OF THE PRODUCT(S) AND RELATED DOCUMENTATION EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.



Mellanox Technologies  
350 Oakmead Parkway Suite 100  
Sunnyvale, CA 94085  
U.S.A.  
[www.mellanox.com](http://www.mellanox.com)  
Tel: (408) 970-3400  
Fax: (408) 970-3403

© Copyright 2019. Mellanox Technologies Ltd. All Rights Reserved.

Mellanox®, Mellanox logo, Connect-IB®, ConnectX®, CORE-Direct®, GPUDirect®, LinkX®, Mellanox Multi-Host®, Mellanox Socket Direct®, UFM®, and Virtual Protocol Interconnect® are registered trademarks of Mellanox Technologies, Ltd.

For the complete and most updated list of Mellanox trademarks, visit <http://www.mellanox.com/page/trademarks>.

All other trademarks are property of their respective owners.

# Table of Contents

<b>Release Update History</b> .....	<b>6</b>
<b>Chapter 1 Overview</b> .....	<b>7</b>
1.1 Supported Devices .....	7
1.2 Supported Cables and Modules .....	9
1.2.1 Validated and Supported 1GbE Cables .....	9
1.2.2 Validated and Supported 10GbE Cables .....	9
1.2.3 Validated and Supported 25GbE Cables .....	11
1.2.4 Validated and Supported 40GbE Cables .....	12
1.2.5 Validated and Supported 50GbE Cables .....	12
1.2.6 Validated and Supported 100GbE Cables .....	12
1.3 Tested Switches .....	14
1.3.1 Tested 10GbE Switches .....	14
1.3.2 Tested 40GbE Switches .....	14
1.3.3 Tested 100GbE Switches .....	14
1.4 Tools, Switch Firmware and Driver Software .....	15
1.5 Supported FlexBoot, UEFI .....	16
1.6 Revision Compatibility .....	16
<b>Chapter 2 Changes and New Features in Rev 14.26.1040</b> .....	<b>17</b>
<b>Chapter 3 Known Issues</b> .....	<b>18</b>
<b>Chapter 4 Bug Fixes History</b> .....	<b>22</b>
<b>Chapter 5 Firmware Changes and New Feature History</b> .....	<b>25</b>
<b>Chapter 6 PreBoot Drivers (FlexBoot/UEFI)</b> .....	<b>35</b>
6.1 FlexBoot Changes and New Features .....	35
6.2 UEFI Changes and Major New Features .....	35
<b>Chapter 7 Unsupported Features and Commands</b> .....	<b>36</b>
7.1 Unsupported Features .....	36
7.2 Unsupported Commands .....	36
<b>Chapter 8 Supported Non-Volatile Configurations</b> .....	<b>37</b>

## List of Tables

Table 1:	Release Update History	6
Table 2:	Supported Devices	7
Table 3:	Validated and Supported 1GbE Cables	9
Table 4:	Validated and Supported 10GbE Cables	9
Table 5:	Validated and Supported 25GbE Cables	11
Table 6:	Validated and Supported 40GbE Cables	12
Table 7:	Validated and Supported 50GbE Cables	12
Table 8:	Validated and Supported 100GbE Cables	12
Table 9:	Tested 10GbE Switches	14
Table 10:	Tested 40GbE Switches	14
Table 11:	Tested 100GbE Switches	14
Table 12:	Tools, Switch Firmware and Driver Software	15
Table 13:	Supported FlexBoot, UEFI	16
Table 14:	Changes and New Features in Rev 14.26.1040	17
Table 15:	Ethernet Rate Limit per VF in RoCE Mode Limitations	18
Table 16:	Ethernet Rate Limit per VF in InfiniBand Mode Limitations	18
Table 17:	Known Issues	18
Table 18:	Bug Fixes History	22
Table 19:	Firmware Changes and New Feature History	25
Table 20:	Unsupported Features	36
Table 21:	Supported Non-Volatile Configurations	37

## Release Update History

*Table 1 - Release Update History*

Release	Date	Description
Rev 14.26.1040	September 29, 2019	Initial version of this firmware release. This version introduces New Features <a href="#">Section 2, “Changes and New Features in Rev 14.26.1040”</a> , on page 17 and Bug Fixes (see <a href="#">Section 4, “Bug Fixes History”</a> , on page 22).

# 1 Overview

These are the release notes for the ConnectX®-4 Lx adapters firmware Rev 14.26.1040.

This firmware supports the following protocols:

- Ethernet - 1GbE, 10GbE, 25GbE, 40GbE, 50GbE,
- PCI Express 3.0, supporting backwards compatibility for v2.0 and v1.1

## 1.1 Supported Devices

This firmware supports the devices and protocols listed in [Table 2](#)

**Table 2 - Supported Devices (Sheet 1 of 3)**

Device Part Number	PSID	Device Name	FlexBoot	UEFI x86	UEFI ARM	Enable/disable exprom Feature
MCX4111A-ACAT	MT_2410110034	ConnectX®-4 Lx EN network interface card, 25GbE single-port SFP28, PCIe3.0 x8, tall bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4111A-XCAT	MT_2410110004	ConnectX®-4 Lx EN network interface card, 10GbE single-port SFP28, PCIe3.0 x8, tall bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4121A-ACAT	MT_2420110034	ConnectX®-4 Lx EN network interface card; 25GbE dual-port SFP28; PCIe3.0 x8; ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4121A-XCAT	MT_2420110004	ConnectX®-4 Lx EN network interface card, 10GbE dual-port SFP28, PCIe3.0 x8, tall bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4131A-BCAT	MT_2430110027	ConnectX®-4 Lx EN network interface card, 40GbE single-port QSFP28, PCIe3.0 x8, tall bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4131A-GCAT	MT_2430110032	ConnectX®-4 Lx EN network interface card, 50GbE single-port QSFP28, PCIe3.0 x8, tall bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4411A-ACAN	MT_2450111034	ConnectX®-4 Lx EN network interface card for OCP; 25GbE single-port SFP28; PCIe3.0 x8; ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4411A-ACQN	MT_2450112034	ConnectX®-4 Lx EN network interface card for OCP with Host Management, 25GbE single-port SFP28, PCIe3.0 x8, no bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4421A-ACAN	MT_2470111034	ConnectX®-4 Lx EN network interface card for OCP, 25GbE dual-port SFP28, PCIe3.0 x8, no bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists

**Table 2 - Supported Devices (Sheet 2 of 3)**

Device Part Number	PSID	Device Name	FlexBoot	UEFI x86	UEFI ARM	Enable/disable exprom Feature
MCX4421A-ACQN	MT_2470112034	ConnectX®-4 Lx EN network interface card for OCP with Host Management, 25GbE dual-port SFP28, PCIe3.0 x8, no bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4421A-ACUN	MT_0000000275	ConnectX®-4 Lx EN network interface card for OCP2.0, Type 1, without host management, 25GbE dual-port SFP28, PCIe3.0 x8, UEFI Enabled, no bracket	Present (Enabled)	Present (Enabled)	Present (Disabled)	Exists
MCX4421A-XCQN	MT_2470110004	ConnectX®-4 Lx EN network interface card for OCP with Host Management, 10GbE dual-port SFP28, PCIe3.0 x8, no bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Not Present	Exists
MCX4431A-GCAN	MT_2490111032	ConnectX®-4 Lx EN network interface card for OCP, with Host Management, 50GbE single-port QSFP28, PCIe3.0 x8, no bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4431M-GCAN	MT_2510111032	ConnectX®-4 Lx EN network interface card for OCP with Multi-Host and Host Management, 50GbE single-port QSFP28, PCIe3.0 x8, no bracket, ROHS R6	Present (Enabled)	Present (Disabled)	Present (Disabled)	Exists
MCX4111A-ACUT	MT_0000000267	ConnectX®-4 Lx EN network interface card, 25GbE single-port SFP28, PCIe3.0 x8, UEFI Enabled, tall bracket	Present (Enabled)	Present (Enabled)	Present (Disabled)	Exists
MCX4121A-ACUT	MT_0000000266	ConnectX®-4 Lx EN network interface card, 25GbE dual-port SFP28, PCIe3.0 x8, UEFI Enabled, tall bracket	Present (Enabled)	Present (Enabled)	Present (Disabled)	Exists
MCX4621A-ACAB	MT_0000000238	ConnectX®-4 Lx EN network interface card for OCP 3.0, with host management, 25GbE Dual-port SFP28, PCIe3.0 x8, Thumbscrew bracket	Present (Enabled)	Present (Enabled)	Present (Disabled)	Exists
MCX4121A-XCHT	MT_0000000414	ConnectX®-4 Lx EN network interface card, with host management, 10GbE dual-port SFP28, PCIe3.0 x8, UEFI Enabled, tall bracket	Present (Enabled)	Present (Enabled)	Present (Disabled)	Exists
MCX4411A-ACUN	MT_0000000268	ConnectX®-4 Lx EN network interface card for OCP, without host management, 25GbE single-port SFP28, PCIe3.0 x8, UEFI Enabled, no bracket	Present (Enabled)	Present (Enabled)	Present (Disabled)	Exists

**Table 2 - Supported Devices (Sheet 3 of 3)**

Device Part Number	PSID	Device Name	FlexBoot	UEFI x86	UEFI ARM	Enable/disable exprom Feature
MCX4431A-GCA	MT_2490110032	ConnectX®-4 Lx EN network interface card for OCP, with Host Management, 50GbE single-port QSFP28, PCIe3.0 x8, no bracket	Present (Enabled)	Not Present	Not Present	Not Available

## 1.2 Supported Cables and Modules

Please refer to the LinkX® Cables and Transceivers web page

(<http://www.mellanox.com/products/interconnect/cables-configurator.php>) for the list of supported cables.

### 1.2.1 Validated and Supported 1GbE Cables

### 1.2.2 Validated and Supported 10GbE Cables

**Table 4 - Validated and Supported 10GbE Cables**

Speed	Cable OPN #	Description
10GbE	1-2053783-1	SFP-H10G-CU1M
10GbE	BN-QS-SP-CBL-5M	40G QSFP+ to 4xSFP+ DAC Breakout Direct Attach Cable 5m
10GbE	BN-QS-SP-CBL-5M	40G QSFP+ to 4xSFP+ DAC Breakout Direct Attach Cable 5m
10GbE	CAB-SFP-SFP-1M	Arista 10GBASE-CR SFP+ Cable 1 Meter
10GbE	CAB-SFP-SFP-1M	Arista Compatible 10G SFP+ Passive Cable 1m
10GbE	CAB-SFP-SFP-3M	Arista 10GBASE-CR SFP+ Cable 3 Meter
10GbE	CAB-SFP-SFP-5M	Arista 10GBASE-CR SFP+ Cable 5 Meter
10GbE	CAB-SFP-SFP-5M	Arista Compatible 10G SFP+ Passive Cable 5m
10GbE	FTLX1471D3BCL-ME	10GBASE-LR SFP+ 1310nm 10km DOM Transceiver Module
10GbE	FTLX8570D3BCL-C2	Cisco FET-10G 10-2566-02 FTLX8570D3BCL-C2 10Gbps Fabric Extender SFP+ Module
10GbE	MC2309124-004	Mellanox Passive Copper Cable ETH 10GBE 10GB/S QSFP TO SFP+ 4M
10GbE	MC2309124-005	Mellanox Passive Copper Cable ETH 10GBE 10GB/S QSFP TO SFP+ 5M
10GbE	MC2309130-001	Mellanox Passive Copper Cable ETH 10GBE 10GB/S QSFP TO SFP+ 1M
10GbE	MC2309130-002	Mellanox Passive Copper Cable ETH 10GBE 10GB/S QSFP TO SFP+ 2M
10GbE	MC2309130-003	Mellanox Passive Copper Cable ETH 10GBE 10GB/S QSFP TO SFP+ 3M
10GbE	MC2309130-00A	Mellanox Passive Copper Cable ETH 10GBE 10GB/S QSFP TO SFP+ 0.5M
10GbE	MC2609125-004	Mellanox Passive Copper Hybrid Cable ETH 40GbE TO 4X10GBE QSFP TO 4X SFP+ 4M
10GbE	MC2609130-001	Mellanox Passive Copper Hybrid Cable ETH 40GbE TO 4X10GBE QSFP TO 4X SFP+ 1M



**Table 4 - Validated and Supported 10GbE Cables**

Speed	Cable OPN #	Description
10GbE	MC2609130-002	Mellanox Passive Copper Hybrid Cable ETH 40GbE TO 4X10GBE QSFP TO 4X SFP+ 2M
10GbE	MC2609130-003	Mellanox Passive Copper Hybrid Cable ETH 40GbE TO 4X10GBE QSFP TO 4X SFP+ 3M
10GbE	MC2609130-0A1	Mellanox Passive Copper Hybrid Cable ETH 40GbE TO 4X10GBE QSFP TO 4X SFP+ 1.5M
10GbE	MC3309124-004	Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 4M
10GbE	MC3309124-005	Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 5M
10GbE	MC3309124-006	Mellanox® Passive Copper Cable, ETH 10GbE, 10Gb/s, SFP+, 6m
10GbE	MC3309124-007	Mellanox® Passive Copper Cable, ETH 10GbE, 10Gb/s, SFP+, 7m
10GbE	MC3309130-001	Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 1M
10GbE	MC3309130-002	Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 2M
10GbE	MC3309130-003	Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 3M
10GbE	MC3309130-004	Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 4M
10GbE	MC3309130-005	Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 5M
10GbE	MC3309130-006	Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 6M
10GbE	MC3309130-007	Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 7M
10GbE	MC3309130-00A	Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 0.5M
10GbE	MC3309130-0A1	Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 1.5M
10GbE	MC3309130-0A2	Mellanox Passive Copper Cable ETH 10GBE 10GB/S SFP+ 2.5M
10GbE	MFMT02A-SR-P	Mellanox Optical Module ETH 10GbE 10GB/S SFP+ LC-LC 850NM SR up to 300M
10GbE	SFP-10G-SR	Cisco 10GBASE-SR SFP+ transceiver module for MMF, 850-nm wavelength, LC duplex connector
10GbE	SFP-H10GB-CU1M	Cisco 1-m 10G SFP+ Twinax cable assembly, passive
10GbE	SFP-H10GB-CU3M	Cisco 3-m 10G SFP+ Twinax cable assembly, passive
10GbE	SFP-H10GB-CU5M	Cisco 5-m 10G SFP+ Twinax cable assembly, passive
10GbE	FTLX8571D3BCL-ME	10gb SFP 850nm Optic Transceiver
10GbE	FTLX1471D3BCL-ME	10GBASE-LR SFP+ 1310nm 10km DOM Transceiver Module
10GbE	74752-9521 (SFP-H10GB-CU5M)	Cisco 10GBASE SFP+ modules
10GbE	74752-9096 (SFP-H10GB-SU5M)	Cisco-Molex INC Active DAC SFP+ 5m

### 1.2.3 Validated and Supported 25GbE Cables



The 25GbE cables can be supported in ConnectX-4 adapter cards only when connected to the MAM1Q00A-QSA28 module.

**Table 5 - Validated and Supported 25GbE Cables**

Speed	Cable OPN #	Description
25GbE	FTLF1436P3BCL	25GE LR 10km SFP28 Optical Transceiver
25GbE	FTLF8536P4BCL	Finisar SFP+ transceivers 25Gb/s
25GbE	FTLF8536P4BCL	25GE SR SFP+ Optical Transceiver
25GbE	LTF8507-PC07	Hisense active fiber cable, 25GbE
25GbE	MCP2M00-A001	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 1m
25GbE	MCP2M00-A002	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 2m
25GbE	MCP2M00-A003	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 3m
25GbE	MCP2M00-A003AP	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 3m, 26AWG
25GbE	MCP2M00-A005E26L	Mellanox Passive Copper Cable, ETH, up to 25GB/S, SFP28, 5M, black, 26AWG, CA-L
25GbE	MCP2M00-A00A	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 0.5m
25GbE	MCP2M00-A01A	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 1.5m
25GbE	MCP2M00-A02A	Mellanox® Passive Copper cable, ETH, up to 25Gb/s, SFP28, 2.5m
25GbE	MCP7F00-A001	Mellanox Passive Copper Hybrid cable ETH 100GbE to 4X25GBS QSFP28 to 4XSFP28 1M
25GbE	MCP7F00-A002	Mellanox Passive Copper Hybrid Cable ETH 100GbE TO 4X25GBS QSFP28 TO 4XSFP28 2M
25GbE	MCP7F00-A003	Mellanox Passive Copper Hybrid Cable ETH 100GbE TO 4X25GBS QSFP28 TO 4XSFP28 3M
25GbE	MCP7F00-A003-AM	Mellanox® passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 3M 30AWG
25GbE	MCP7F00-A005AM	Mellanox Passive Copper Hybrid Cable ETH 100GbE TO 4X25GBS QSFP28 to 4XSFP28 5M
25GbE	MCP7F00-A01A	Mellanox Passive Copper Hybrid Cable ETH 100GbE to 4X25GBS QSFP28 to 4XSFP28 1.5M
25GbE	MCP7F00-A02A	Mellanox Passive Copper Hybrid Cable ETH 100GbE to 4X25GBS QSFP28 to 4XSFP28 2.5M
25GbE	MFA2P10-Axxx	Mellanox® active optical cable 25GbE, SFP28, up to 50m
25GbE	MMA2P00-AS	Mellanox® transceiver, 25GbE, SFP28, LC-LC, 850nm, SR, up to 100m
25GbE	MFM1T02A-SR-P	Mellanox® Optical Module ETH 10GbE 10GB/S SFP+ LC-LC 850NM SR up to 300M
25GbE	SFP-H25G-CU1M	25GBASE-CR1 Copper Cable 1-meter
25GbE	SFP-H25G-CU2M	25GBASE-CR1 Copper Cable 2-meter
25GbE	SFP-H25G-CU3M	25GBASE-CR1 Copper Cable 3-meter

## 1.2.4 Validated and Supported 40GbE Cables

**Table 6 - Validated and Supported 40GbE Cables**

Speed	Cable OPN #	Description
40GbE	00D5811-N13445C	IBM Cable DAC 40GbE QSFP+ to QSFP+ Passive Copper 7m
NA	MAM1Q00A-QSA	Mellanox® cable module, ETH 10GbE, 40Gb/s to 10Gb/s, QSFP to SFP+
NA	MAM1Q00A-QSA28	Mellanox® cable module, ETH 25GbE, 100Gb/s to 25Gb/s, QSFP28 to SFP28
40GbE	MC2210126-004	Mellanox® Passive Copper Cable, ETH 40GbE, 40GbE, QSFP, 4m
40GbE	MC2210126-005	Mellanox® Passive Copper Cable, ETH 40GbE, 40GbE, QSFP, 5m
40GbE	MC2210128-003	Mellanox Passive Copper Cable ETH 40GbE 40GbE QSFP 3M
40GbE	MC2210130-001	Mellanox Passive Copper Cable ETH 40GbE 40GbE QSFP 1M
40GbE	MC2210130-002	Mellanox Passive Copper Cable ETH 40GbE 40GbE QSFP 2M
40GbE	MC2210130-00A	Mellanox® Passive Copper Cable, ETH 40GbE, 40GbE, QSFP, 0.5m
40GbE	MC2210130-00B	Mellanox® Passive Copper Cable, ETH 40GbE, 40GbE, QSFP, 0.75m
40GbE	MC2210411-SR4	Mellanox Optical Module 40GbE QSFP MPO 850NM UP TO 100M
40GbE	MC2210411-SR4E	Mellanox Optical Module 40GbE QSFP MPO 850NM UP TO 300M
40GbE	QSFP-40G-SR-BD	Cisco 40GBASE-SR-BiDi, duplex MMF
40GbE	QSFP-40G-SR4	Cisco 40GBASE-SR4, 4 lanes, 850 nm MMF
40GbE	QSFP-H40G-ACU10M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 10-meter, active
40GbE	QSFP-H40G-AOC10M	Cisco 40GBase-AOC QSFP direct-attach Active Optical Cable, 10-meter
40GbE	QSFP-H40G-CU1M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 1-meter, passive
40GbE	QSFP-H40G-CU3M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 3-meter, passive
40GbE	QSFP-H40G-CU5M	Cisco 40GBASE-CR4 QSFP direct-attach copper cable, 5-meter, passive

## 1.2.5 Validated and Supported 50GbE Cables

**Table 7 - Validated and Supported 50GbE Cables**

Speed	Cable OPN #	Description
50GbE	MCP7H00-G01A	Mellanox Passive Copper Hybrid Cable ETH 100GbE TO 2X50GBS QSFP28 TO 2XQSFP28 1.5M
50GbE	MCP7H00-G02A	Mellanox Passive Copper Hybrid Cable ETH 100GbE TO 2X50GBS QSFP28 TO 2XQSFP28 2.5M

## 1.2.6 Validated and Supported 100GbE Cables

**Table 8 - Validated and Supported 100GbE Cables**

Speed	Cable OPN #	Description
100GbE	CBL-00195-02	100GbE QSFP28 to QSFP28 copper cable 3M
100GbE	MCP1600-C001	Mellanox Passive Copper Cable ETH 100GbE 100GBS QSFP LSZH 1M
100GbE	MCP1600-C002	Mellanox Passive Copper Cable ETH 100GbE 100GBS QSFP LSZH 2M

**Table 8 - Validated and Supported 100GbE Cables**

Speed	Cable OPN #	Description
100GbE	MCP1600-C003	Mellanox Passive Copper Cable ETH 100GbE 100GBS QSFP LSZH 3M
100GbE	MCP1600-C005AM	Mellanox® Passive Copper cable, ETH 100GbE, 100GbE, QSFP, 5m, 26AWG
100GbE	MCP1600-C005E26L	Mellanox® Passive Copper cable, ETH 100GbE, 100GbE, QSFP28, 5m, Black, 26AWG, CA-L
100GbE	MCP1600-C00A	Mellanox Passive Copper Cable ETH 100GbE 100GBS QSFP LSZH 0.5M
100GbE	MCP1600-C01A	Mellanox® Passive Copper cable, ETH 100GbE, 100GbE, QSFP, LSZH, 1.5m
100GbE	MCP1600-C02A	Mellanox® Passive Copper cable, ETH 100GbE, 100GbE, QSFP, LSZH, 2.5m
100GbE	MCP1600-C03A	Mellanox® Passive Copper cable, ETH 100GbE, 100GbE, QSFP, PVC, 3.5m 26AWG
100GbE	MCP7F00-A005R26L	Mellanox® passive copper hybrid cable, ETH 100GbE to 4x25GbE, QSFP28 to 4xSFP28, 5m, Colored, 26AWG, CA-L
100GbE	MFA1A00-C003	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 3m
100GbE	MFA1A00-C005	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 5m
100GbE	MFA1A00-C010	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 10m
100GbE	MFA1A00-C015	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 15m
100GbE	MFA1A00-C020	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 20m
100GbE	MFA1A00-C030	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 30m
100GbE	MFA1A00-C050	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 50m
100GbE	MFA1A00-C100	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 100m
100GbE	MFS1200-C005	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 5m
100GbE	MFS1200-C010	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 10m
100GbE	MFS1200-C015	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 15m
100GbE	MFS1200-C020	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 20m
100GbE	MFS1200-C030	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 30m
100GbE	MFS1200-C050	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 50m
100GbE	MFS1200-C100	Mellanox® Active Fiber Cable, ETH 100GbE, 100GbE, QSFP, LSZH, 100m
100GbE	MMA1B00-C100_B	Mellanox® transceiver, up to 100GbE, QSFP28, MPO, 850nm, up to 100m OM3
100GbE	MMA1B00-C100D	Mellanox® Transceiver, 100GbE, QSFP28, MPO, 850nm, up to 100m
100GbE	MMS1C00-C500	Mellanox® transceiver, 100GbE, QSFP28, MPO, 1550nm PSM4, up to 2km
100GbE	MMS1C00-C500	Mellanox® transceiver, 100GbE, QSFP28, MPO, 1550nm PSM4, up to 2km

## 1.3 Tested Switches

### 1.3.1 Tested 10GbE Switches

**Table 9 - Tested 10GbE Switches**

Speed	Switch Silicon	OPN # / Name	Description	Vendor
10/40GbE	N/A	3064	48-port 10Gb/40Gb Switch	Cisco
10/40GbE	N/A	7050Q	16-port 40Gb Switch	Arista
10/40GbE	N/A	7050S	48-port 10Gb/40Gb Switch	Arista
10GbE	N/A	5548	Cisco 10GB ETH switch	Cisco
10GbE	N/A	G8264	BNT 10/40GB ETH switch	BNT
10GbE	N/A	QFX3500	Juniper 10/40GB ETH switch	Juniper
10GbE	N/A	S4810P-AC	48-port 10Gb/40Gb Switch	Force10
10GbE	SwitchX®	SX1016X-1BFR	64-Port 10GbE Switch System	Mellanox

### 1.3.2 Tested 40GbE Switches

**Table 10 - Tested 40GbE Switches**

Speed	Switch Silicon	OPN # / Name	Description	Vendor
10/40GbE	N/A	3064	48-port 10Gb/40Gb Switch	Cisco
10/40GbE	N/A	7050Q	16-port 40Gb Switch	Arista
10/40GbE	N/A	7050S	48-port 10Gb/40Gb Switch	Arista
40GbE	N/A	3132Q	Cisco 40GB ETH switch	Cisco
40GbE	N/A	7050QX	32-port 40Gb Switch	Arista
40GbE	N/A	G8316	BNT 40GB RackSwitch G8316	BNT
40GbE	N/A	S6000	32-port 40Gb Switch	Dell
40GbE	SwitchX®	SX1036B-1BFR	36-Port 40/56GbE Switch System	Mellanox

### 1.3.3 Tested 100GbE Switches

**Table 11 - Tested 100GbE Switches**

Speed	Switch Silicon	OPN # / Name	Description	Vendor
100GbE	N/A	7060CX	32-port 100Gb Switch	Arista
100GbE	N/A	93180YC-EX	48 x 10/25-Gbps fiber ports and 6 x 40/100-Gbps Quad Small Form-Factor Pluggable 28 (QSFP28) ports	Cisco
100GbE	N/A	C3232C	High-Density, 100 Gigabit Ethernet Switch	Cisco
100GbE	N/A	CE8860-4C-EI	24x10GE (SFP+) or 25GE (SFP28) and 2x100GE switch	Huawei

**Table 11 - Tested 100GbE Switches**

Speed	Switch Silicon	OPN # / Name	Description	Vendor
100GbE	Spectrum	SN2410-CB2F	48-port 25GbE + 8-port 100GbE Open Ethernet ToR Switch System	Mellanox
100GbE	Spectrum	SN2700-CS2R	32-port Non-blocking 100GbE Open Ethernet Spine Switch System	Mellanox
100GbE	Spectrum	SN2740-CB2F1	32-port Non-blocking 100GbE Open Ethernet Spine Switch System	Mellanox
100GbE	N/A	Wedge 100-32X R04	32-port 100GbE QSFP28 - Leaf/Spine Switch, power-to-port airflow, DC Power	Edgecore
100GbE	N/A	QFX5200-32C-32	32-port 100GbE Ethernet Switch System	Juniper
100GbE	N/A	S6820-56HF	48 SFP+ + 8 QSFP Ports 100GbE Switch Ethernet	H3C

## 1.4 Tools, Switch Firmware and Driver Software

The following are the drivers' software, tools, switch/HCA firmware versions tested that you can upgrade from or downgrade to when using firmware Rev 14.26.1040:

**Table 12 - Tools, Switch Firmware and Driver Software**

	Supported Version
MLNX_OFED	4.7-x.0.0.0 / 4.6-1.0.1.1
MLNX_EN (MLNX_OFED based code)	4.7-x.0.0.0 / 4.6-1.0.1.1
WinOF-2	2.30 / 2.20
MFT	4.13.0 / 4.12.0
MLNX-OS	3.8.2004
Onyx	3.8.2004
ConnectX-4 Lx Firmware	14.25.1020 / 14.24.1000
Linux Inbox Drivers	<ul style="list-style-type: none"> <li>• RH7.6</li> <li>• Ubuntu 16.04.05</li> </ul>
Windows Inbox Drivers	<ul style="list-style-type: none"> <li>• Windows 2012</li> <li>• Windows 2012 R2</li> <li>• Windows 2016</li> </ul>

## 1.5 Supported FlexBoot, UEFI



Please be aware that not all firmware binaries contain FlexBoot or UEFI, support may vary between cards (see [Section 1.1, “Supported Devices”, on page 7.](#))

Firmware Rev 14.26.1040 supports the following FlexBoot:

**Table 13 - Supported FlexBoot, UEFI**

Expansion ROM	Supported Version
FlexBoot	3.5.803
UEFI	14.19.14

## 1.6 Revision Compatibility

Firmware Rev 14.26.1040 complies with the following programmer’s reference manual:

- *Mellanox Adapters Programmer’s Reference Manual (PRM), Rev 0.47 or later*, which has Command Interface Revision 0x5. The command interface revision can be retrieved by means of the QUERY\_FW command and is indicated by the field `cmd_interface_rev`.

## 2 Changes and New Features in Rev 14.26.1040

**Table 14 - Changes and New Features in Rev 14.26.1040**

Feature/Change	Description
<b>Rev. 14.26.1040</b>	
<b>ICMD and Diagnostic Counters</b>	Enabled the firmware by using the ICMD commands to deal with diagnostic counters similar to cmdif. They can be called via the vsec space. The counters' values are returned only via the tracer. The ICMD Query Caps indicate support and expose the list of the supported counters.
<b>User Context Object (DEVX)</b>	This is a containerized sandbox per user, to access PRM command securely by using General Object commands, UMEM and UCTX contexts. The allowed functionalities of this capability depend on the user permissions. The following functionalities are still managed by the Kernel: <ul style="list-style-type: none"> <li>• Resource cleaning</li> <li>• UCTX stamping</li> <li>• Blocking the physical address and IRQ from these UCTX</li> </ul>
<b>DEVX Support for Asynchronous Events</b>	Added support for reporting the supported affiliated and unaffiliated asynchronous events to DEVX users through the command interface.
<b>Zero-Touch-RoCE Counters</b>	Zero-Touch-RoCE counters are now available to the user for debuggability purposes when using the Zero-Touch-RoCE feature.
<b>Security Hardening Enhancements</b>	This release contains important reliability improvements and security hardening enhancements. Mellanox recommends upgrading your device firmware to this release to improve the device firmware security and reliability.
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”, on page 22</a>



### 3 Known Issues

The following table describes known issues in this firmware release and possible workarounds.

For a list of old firmware Know Issues, please see ConnectX-Lx Firmware Archived Known Issues file ([http://www.mellanox.com/pdf/firmware/ConnectX4Lx-Firmware\\_Archived\\_Known\\_Issues.pdf](http://www.mellanox.com/pdf/firmware/ConnectX4Lx-Firmware_Archived_Known_Issues.pdf))

**Table 15 - Ethernet Rate Limit per VF in RoCE Mode Limitations**

Adapter Card	Dual Port Device				Single Port Device	
	w/o LAG (TOTAL_VFS>32)		With LAG (TOTAL_VFS<32)		w/o LAG	
	w/o QoS	Full QoS	w/o QoS	Full QoS	w/o QoS	Full QoS
ConnectX-4 Lx	127	45	32	20	127	100

**Table 16 - Ethernet Rate Limit per VF in InfiniBand Mode Limitations**

Adapter Card	Dual Port Device		Single Port Device	
	w/o LAG		w/o LAG	
	w/o QoS	Full QoS	w/o QoS	Full QoS
ConnectX-4 Lx	127	26	127	55

**Table 17 - Known Issues (Sheet 1 of 4)**

Internal Ref.	Issue
1919403	<b>Description:</b> Hardware arbitration is currently disabled in OCP3.0 cards. It will be supported on future releases for the same hardware.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Hardware arbitration, OCP3.0
	<b>Discovered in Version:</b> 14.26.1040
1840289	<b>Description:</b> Since Packet Pacing enforce max_tc value is "1", features that require multiple TCs will not be active when this mode is available.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Packet Pacing
	<b>Discovered in Version:</b> 14.26.1040
1796628	<b>Description:</b> Due to performance considerations, unicast loopback traffic will go through the NIC SX tables, and multicast loopback traffic will skip the NIC SX tables.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Performance, unicast loopback traffic, multicast loopback traffic
	<b>Discovered in Version:</b> 14.26.1040
1699214	<b>Description:</b> NODNIC VF is partially tested. It is fully tested only in ConnectX-5 adapter cards.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> NODNIC VF
	<b>Discovered in Version:</b> 14.25.1020

**Table 17 - Known Issues (Sheet 2 of 4)**

Internal Ref.	Issue
1689186	<p><b>Description:</b> Changing priority to TC map during traffic might cause packet drops.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> QoS</p> <p><b>Discovered in Version:</b> 14.25.1020</p>
1604699	<p><b>Description:</b> Ethernet RFC 2819 counter ether_stats_oversize_pkts and Ethernet IEEE 802.3 counter a_frame_too_long_errors share the same resource. Clearing each of them will affect the other.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Counters</p> <p><b>Discovered in Version:</b> 14.25.1020</p>
-	<p><b>Description:</b> In Ethernet mode, at 10/40GbE speeds, only NO-FEC in Force mode is supported. Other user configurations are overridden.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Ethernet, 10GbE, 40GbE, RS-FEC</p> <p><b>Discovered in Version:</b> 14.25.1020</p>
1498399	<p><b>Description:</b> If the XRC switches between SRQ/RMPs while there is an outstanding ODP on the responder XRC QP, a CQE with an error might be generated (that is not a PFAULT abort).</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> XRC SRQ/RMP ODP</p> <p><b>Discovered in Version:</b> 14.25.1020</p>
1426283	<p><b>Description:</b> An mlxconfig configuration followed by the init 0 command might cause the HCA to remain in high power instead of entering the standby state. Such behavior will require a server reboot after running the mlxconfig tool.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> mlxconfig</p> <p><b>Discovered in Version:</b> 14.24.1000</p>
1546401	<p><b>Description:</b> vport_tc and para_vport_tc are not supported in this version.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> SR-IOV vport_tc and para_vport_tc</p> <p><b>Discovered in Version:</b> 14.24.1000</p>
1566930	<p><b>Description:</b> In case a faulty cable is inserted, (either power issue or I2C error), the traffic LED is turned OFF although it is expected to be ON. and the amber LED is blinking as expected.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Cables</p> <p><b>Discovered in Version:</b> 14.24.1000</p>
1546492	<p><b>Description:</b> Executing the update_lid command while the IB port sniffer utility is active can stop the utility.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> IB Sniffer</p> <p><b>Discovered in Version:</b> 14.24.1000</p>

**Table 17 - Known Issues (Sheet 3 of 4)**

Internal Ref.	Issue																																
1537898	<b>Description:</b> Initializing a function while the IB port sniffer utility is active can stop the utility.																																
	<b>Workaround:</b> N/A																																
	<b>Keywords:</b> IB Sniffer																																
	<b>Discovered in Version:</b> 14.24.1000																																
1332714	<b>Description:</b> The maximum “read” size of MTRC_STDB is limited to 272 Bytes.																																
	<b>Workaround:</b> Set the MTRC_STDB.read_size to the maximum value of 0x110=272 Bytes																																
	<b>Keywords:</b> Access register, MTRC_STDB, tracer to dmesg, fwtrace to dmesg																																
	<b>Discovered in Version:</b> 14.23.1020																																
1408994	<b>Description:</b> FTE with both forward (FWD) and encapsulation (ENCAP) actions is not supported in the SX NIC Flow Table.																																
	<b>Workaround:</b> N/A																																
	<b>Keywords:</b> SX NIC Flow Table																																
	<b>Discovered in Version:</b> 14.23.1020																																
1350794	<b>Description:</b> Encapsulation / Decapsulation support in steering has the following limitations: <ul style="list-style-type: none"> <li>• Encapsulation / Decapsulation can be open on the FDB only if all VFs are non active.</li> <li>• Encapsulation / Decapsulation supports single mode only: FDB / NIC. Opening tables of both types is not supported.</li> <li>• RoCE is supported only when no Encapsulation / Decapsulation is opened on the FDB.</li> <li>• Encapsulation / Decapsulation per device support: <table border="1" data-bbox="510 976 981 1176"> <thead> <tr> <th></th> <th></th> <th>NIC</th> <th>FDB</th> <th></th> </tr> </thead> <tbody> <tr> <td rowspan="2">ConnectX-4</td> <td>encap</td> <td>NO</td> <td>YES</td> <td>non MH</td> </tr> <tr> <td>decap</td> <td>NO</td> <td>NO</td> <td></td> </tr> <tr> <td rowspan="2">ConnectX-4 Lx</td> <td>encap</td> <td>NO</td> <td>YES</td> <td>non MH</td> </tr> <tr> <td>decap</td> <td>NO</td> <td>YES</td> <td></td> </tr> <tr> <td rowspan="2">ConnectX-5</td> <td>encap</td> <td>YES</td> <td>YES</td> <td></td> </tr> <tr> <td>decap</td> <td>YES</td> <td>YES</td> <td></td> </tr> </tbody> </table> </li> </ul>			NIC	FDB		ConnectX-4	encap	NO	YES	non MH	decap	NO	NO		ConnectX-4 Lx	encap	NO	YES	non MH	decap	NO	YES		ConnectX-5	encap	YES	YES		decap	YES	YES	
			NIC	FDB																													
	ConnectX-4	encap	NO	YES	non MH																												
		decap	NO	NO																													
	ConnectX-4 Lx	encap	NO	YES	non MH																												
decap		NO	YES																														
ConnectX-5	encap	YES	YES																														
	decap	YES	YES																														
<b>Workaround:</b> N/A																																	
<b>Keywords:</b> Steering Encapsulation / Decapsulation																																	
<b>Discovered in Version:</b> 14.23.1020																																	
1027553	<b>Description:</b> While using e-switch vport sVLAN stripping, the RX steering values on the sVLAN might not be accurate.																																
	<b>Workaround:</b> N/A																																
	<b>Keywords:</b> e-sw vport sVLAN stripping, RX steering																																
	<b>Discovered in Version:</b> 14.24.1000																																
1799917	<b>Description:</b> Untagged CVLAN packets in the Steering Flow Tables do not match the SVLAN tagged packets.																																
	<b>Workaround:</b> N/A																																
	<b>Keywords:</b> Steering Flow Tables, CVLAN/SVLAN packets																																
	<b>Discovered in Version:</b> 14.23.1020																																
1355883	<b>Description:</b> Running the QUERY_VPORT_COUNTER command with clear bit results in discard counters being reset.																																
	<b>Workaround:</b> N/A																																
	<b>Keywords:</b> Discard counters																																
	<b>Discovered in Version:</b> 14.22.1002																																

**Table 17 - Known Issues (Sheet 4 of 4)**

Internal Ref.	Issue
1277762	<b>Description:</b> An Ethernet multicast loopback packet is not counted (even if it is not a local loopback packet) when running the <code>nic_receive_steering_discard</code> command.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Ethernet multicast loopback packet
	<b>Discovered in Version:</b> 14.22.1002
1047184	<b>Description:</b> RDMA <code>resq_local_length_error</code> and <code>resp_remote_invalid_request</code> counters do not function properly.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> RDMA counters
	<b>Discovered in Version:</b> 14.21.1000
1168594	<b>Description:</b> RoCE Dual Port Mode (a.k.a Multi-Port vHCA: MPV) is not supported in Multi-Host setups.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Multi-Port vHCA, Multi-Host
	<b>Discovered in Version:</b> 14.21.1000

## 4 Bug Fixes History

Table 18 lists the bugs fixed in this release. For a list of old firmware Bug Fixes, please see ConnectX-4 Lx Firmware Archived Bug Fixes file

([http://www.mellanox.com/pdf/firmware/ConnectX4Lx-Firmware\\_Archived\\_Bug\\_Fixes.pdf](http://www.mellanox.com/pdf/firmware/ConnectX4Lx-Firmware_Archived_Bug_Fixes.pdf))

**Table 18 - Bug Fixes History (Sheet 1 of 3)**

Internal Ref.	Issue
1822787	<b>Description:</b> Fixed an issue that caused a function to misbehave when a PCIe TLP was set with a poisoned indication.
	<b>Keywords:</b> PCIe TLP
	<b>Discovered in Version:</b> 14.25.1020
	<b>Fixed in Release:</b> 14.26.1040
1801266	<b>Description:</b> Fixed a rare issue that resulted in “destroy mkey “ command getting stuck when rebooting the hypervisor.
	<b>Keywords:</b> FLR
	<b>Discovered in Version:</b> 14.25.1020
	<b>Fixed in Release:</b> 14.26.1040
1753801	<b>Description:</b> In this release, the total firmware reset time is increased by 1 second.
	<b>Keywords:</b> Firmware reset
	<b>Discovered in Version:</b> 14.24.1020
	<b>Fixed in Release:</b> 14.25.6000
1771921	<b>Description:</b> Fixed an issue that prevented users with non-port owner privilege from using the “read DCBX access registry key” REGID_DCBX_APP/REGID_DCBX_PARAM.
	<b>Keywords:</b> DCBX
	<b>Discovered in Version:</b> 14.24.1000
	<b>Fixed in Release:</b> 14.25.1020
1615586	<b>Description:</b> Fixed a rare issue that caused the QP to falsely transition into the error state as a result of handling duplicate read/atomic request followed by memory key invalidation.
	<b>Keywords:</b> CQE
	<b>Discovered in Version:</b> 14.24.1000
	<b>Fixed in Release:</b> 14.25.1020
1678824	<b>Description:</b> Fixed an issue that prevented the user to enable the port after disabling it in the VF NODNIC.
	<b>Keywords:</b> VF NODNIC
	<b>Discovered in Version:</b> 14.24.1000
	<b>Fixed in Release:</b> 14.25.1020
1606289	<b>Description:</b> Enlarged the number of modify fields to 16 to avoid IPv6 header rewrite failure.
	<b>Keywords:</b> IPv6 header rewrite
	<b>Discovered in Version:</b> 14.24.1000
	<b>Fixed in Release:</b> 14.25.1020
1627973	<b>Description:</b> Fixed an issue that prevented IB QP counters for Acks/Responses from working as a results the NACK/OOS counters showed as zero.
	<b>Keywords:</b> IB QP counters for Acks/Responses
	<b>Discovered in Version:</b> 14.24.1000
	<b>Fixed in Release:</b> 14.25.1020

**Table 18 - Bug Fixes History (Sheet 2 of 3)**

Internal Ref.	Issue
1554104	<b>Description:</b> Set the stateless offloads cap to be permanently '1'.
	<b>Keywords:</b> Stateless offloads cap
	<b>Discovered in Version:</b> 14.24.1000
	<b>Fixed in Release:</b> 14.25.1020
1547318	<b>Description:</b> Fixed an issue that prevented the system from counting multicast/broadcast traffic on the ETH unicast vport counter when the driver did not specify the MAC address in the FTE match criteria of the Flow Table Entry in the eswitch's FDB table.
	<b>Keywords:</b> Multicast/broadcast traffic
	<b>Discovered in Version:</b> 14.23.1020
	<b>Fixed in Release:</b> 14.24.1000
1447157	<b>Description:</b> Fixed a standby deadlock that was caused when the PCIe reset arrived during traffic.
	<b>Keywords:</b> PCIe
	<b>Discovered in Version:</b> 14.23.1020
	<b>Fixed in Release:</b> 14.24.1000
1284452/ 1282926	<b>Description:</b> Fixed an issue that caused the mlxconfig tool to present all possible expansion ROM images, instead of presenting only the existing images.
	<b>Keywords:</b> mlxconfig
	<b>Discovered in Version:</b> 14.22.1002
	<b>Fixed in Release:</b> 14.24.1000
1487250	<b>Description:</b> Fixed an issue that prevented the SX error handling mechanism from clearing the counters when the max WQE's re-read attempts were completed successfully.
	<b>Keywords:</b> SX error handling mechanism, WQE, counters
	<b>Discovered in Version:</b> 14..23.1020
	<b>Fixed in Release:</b> 14.24.1000
1371032	<b>Description:</b> Fixed an issue related to the VF's EQ error that caused the health buffer to report to the host (PF) when the VF released the EQ buffer's memory before the EQ was destroyed.
	<b>Keywords:</b> EQ, health buffer
	<b>Discovered in Version:</b> 14.20.1010
	<b>Fixed in Release:</b> 14..24.1000
1424873	<b>Description:</b> Modifying VMQoS rate limiter parameters during traffic might cause transmission failure.
	<b>Keywords:</b> VMQoS, rate limiter
	<b>Discovered in Version:</b> 14.22.1002
	<b>Fixed in Release:</b> 14..24.1000
1475993	<b>Description:</b> Aligned the default tuning type in PHY TEST MODE to the device protocol.
	<b>Keywords:</b> PHY
	<b>Discovered in Version:</b> 14.23.1020
	<b>Fixed in Release:</b> 14..24.1000

**Table 18 - Bug Fixes History (Sheet 3 of 3)**

Internal Ref.	Issue
1403211	<p><b>Description:</b> When a device is operating in Safe Mode state, and the user issues the mlxfwreset command, the device might fail to come-up correctly after the reset.  <b>Note:</b> Do not run mlxfwreset when operating in a Safe Mode state.</p>
	<p><b>Keywords:</b> mlxfwreset</p>
	<p><b>Discovered in Version:</b> 14.23.1020</p>
	<p><b>Fixed in Release:</b> 14.24.1000</p>
1431772	<p><b>Description:</b> Fixed an issue that caused the max_qp_retry_freq_exceeded counter (including a CQE with error syndrome 0x97, and the QP moving to error state) to be activated only after exceeding the NIC Vport context max_qp_retry_limit, and not when reaching it.</p>
	<p><b>Keywords:</b> max_qp_retry_freq_exceeded</p>
	<p><b>Discovered in Version:</b> 14.22.1002</p>
	<p><b>Fixed in Release:</b> 14.24.1000</p>
1295606	<p><b>Description:</b> Fixed an issue related to PCIe "Surprise link down" event reporting capability.</p>
	<p><b>Keywords:</b> PCIe</p>
	<p><b>Discovered in Version:</b> 14.22.1002</p>
	<p><b>Fixed in Release:</b> 14.24.1000</p>
1434863	<p><b>Description:</b> Fixed an issue that resulted in the link partner experiencing false active linkup when plugging in a base-T cable to a closed port.</p>
	<p><b>Keywords:</b> Interfaces</p>
	<p><b>Discovered in Version:</b> 14.22.1002</p>
	<p><b>Fixed in Release:</b> 14.24.1000</p>
1424873	<p><b>Description:</b> Modifying VMQoS rate limiter parameters during traffic might cause transmission failure.</p>
	<p><b>Keywords:</b> VMQoS, rate limiter</p>
	<p><b>Discovered in Version:</b> 14.22.1002</p>
	<p><b>Fixed in Release:</b> 14.24.1000</p>

## 5 Firmware Changes and New Feature History

**Table 19 - Firmware Changes and New Feature History (Sheet 1 of 10)**

Feature/Change	Description
<b>Rev. 14.25.1020</b>	
<b>VSC Security</b>	VSC security includes the mechanisms which will prevent a reasonable host from affecting other hosts from using VSC.
<b>Zero Touch RoCE</b>	Zero touch RoCE enables RoCE to operate on fabrics where no PFC nor ECN are configured. This makes RoCE configuration a breeze while still maintaining its superior high performance.
<b>ODP support for SRQ &amp; XRC</b>	Added support for send opcode operations targeting a SRQ/RMP with the receive WQEs using ODP memory. In case the receive WQE receives an ODP, the device will generate ODP notifications (EQE) and PFAULT will abort CQEs. <b>Note:</b> It is recommended to prefetch the memory used by the receive WQEs to reduce ODP occurrence as these have significant latencies and will cause a performance degradation.
<b>Firmware Burning using DMA Pages</b>	This new capability accelerates the firmware burning process by using Direct Memory Access (DMA) pages.
<b>Auto-Sensing when using 25/10GbE Optical Modules</b>	This new capability accelerates the network to auto-sense the port speed and use it when using a 25/10GbE optical module. Meaning, if the used module is 25GbE but the port is a 10GbE port, the speed used for that network will be 10GbE.
<b>Package ID</b>	Enabled Package ID configuration using server strap according OCP 3.0.
<b>DPDK UIO</b>	This capability provides a solution for improving user space drivers development, generic user space IO device services.
<b>mlxconfig</b>	Renamed the BOOT_RETRY_CNT1 parameter to BOOT_RETRY_CNT.
<b>Reduced Firmware Upgrade Time</b>	Reduced firmware upgrade time using mlxfwreset tool to ~3 seconds. Using this capability requires enabling PARTIAL_RESET_EN in mlxconfig and using MFT version 4.12.0 and up. The “PARTIAL” refers to not resetting the port modules (which is not mandatory for firmware upgrades). <b>Note:</b> Currently this capability only supports firmware upgrade and downgrades to firmware versions newer than XX.25.1020.
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”, on page 22</a>
<b>Rev. 14.24.1000</b>	
<b>Layer 3 Encapsulation</b>	Added support for an additional layer (Layer 3) of packet processing at the hypervisor level that enables adding and removing protocol headers (e.g., the MAC address is removed during encapsulation, and added during decapsulation) for the encapsulated traffic.
<b>e-switch Steering Rule</b>	Enabled e-switch steering rule in the NIC without matching it with the Directional MACs (DMAC) protocol. Now the rule is only according to the MC/UC bit.
<b>IB Sniffer Tool</b>	The IB Sniffer utility provides the user the ability to capture the e-switch traffic directly to a hypervisor queue.
<b>Transmission Histogram Counters</b>	Added support for the transmission histogram counter set as part of the Ethernet extended group counters.
<b>Events Generation by the Hardware upon Counter Incrementation</b>	Enabled the hardware to generate an event upon counter incrementation, in order to reduce an overhead from the software from reading rarely updated counters such as error counters.
<b>NODNIC Connectivity</b>	Enables NOIDNIC connectivity to the network through the e-switch and not directly to the physical port.



**Table 19 - Firmware Changes and New Feature History (Sheet 2 of 10)**

Feature/Change	Description
<b>QP and Mkey Values</b>	Enabled setting the QP and the Mkey values by the software upon these resources creation.
<b>PCIe Atomic</b>	Enabled advanced PCIe atomic operations. The HCA will perform PCIe atomic operations as a requestor towards the host memory when receiving compatible atomic messages from the network, and according to the configuration of NV_SW_OFFLOAD_CONFIG pci_atomic_mode field and the PCI AtomicOp Requester Enable bit in the Device Control 2 register.
<b>TIR Destination from the FDB</b>	Enabled a single TIR destination from the FDB.
<b>WRED</b>	Changed the WRED default mode to OFF for Multi-Host adapter cards.
<b>TX Steering Rule on in WQE Ethernet Segment</b>	Added support for TX steering rule on flow_table_metadata in WQE Ethernet segment.
<b>L3 Encapsulation/Decapsulation in the Reformat Context Allocation</b>	Added L3 encapsulation/decapsulation support in the reformat context allocation. <ul style="list-style-type: none"> <li>L3 encapsulation removes L2 headers and adds generic L3 tunnel encapsulation.</li> <li>L3 decapsulation removes the generic L3 tunnel decapsulation and L2 header.</li> </ul>
<b>Flow Steering Header Modification</b>	Added support for flow steering header modification (header rewrite) for IPv4 TTL header for loopback traffic (VF-VF/VF-PF). <b>Note:</b> TTL modification for traffic from the network is currently not supported.
<b>Teardown: Fast Mode</b>	<b>[Developers only]</b> Moved the fast teardown HCA cap bit to offset 0x1c.4:1.
<b>Virtual Functions/QoS</b>	Enabled Virtual Functions to read QPDPM/QPDP/QPTS.
<b>Message vs. Payload based flow control QP Configuration</b>	Added support for requester QP packet based on E2E credits mode. The new flow control supports HCA-to-switch RDMA traffic packet-based End-2-End.
<b>Multi PCI RDMA IB</b>	This capability enables the user to expose two PCI/IB devices per network port.
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”, on page 22</a>
<b>Steering</b>	Enabled a single TIR destination from the FDB.
<b>Rev. 14.23.1020</b>	
<b>Virtual Functions (VF)</b>	Increased the number of VFs that can work with full VMQoS (8 TC) per PFs as follow: <ul style="list-style-type: none"> <li>in dual port devices to: 0-21, 33-45 VFs (22-32 VFs has single TC)</li> <li>in single port devices to: 0-64 VFs</li> </ul>
<b>InfiniBand</b>	Added support for IPoIB non-default Partition Keys (PKeys). Now the PKey values can be modified in the PKey table without the need of recreating the IPoIB (underlay) QPs.
<b>SR-IOV in Multi-Host/Socket-Direct</b>	<b>[Beta]</b> Added support for SR-IOV (up to 63 VFs) in Multi-Host/Socket-Direct.
<b>Virtualization</b>	Reduced firmware’s memory consumption to increase the supported number of VFs per PF to up to 100.
<b>Tools/Driver Version</b>	Added support for QUERY_DRIVER_VERSION command. This command allows the PF driver to query its VFs driver version which was set by the SET_DRIVER_VERSION command.
<b>Resiliency</b>	Shutting Down RDMA QPs with Excessive Retransmissions is a mechanism used to detect excessive retransmissions for an RC connection, and to close the connection in response to it. If the number of retransmissions due to a Local Ack Timeout, NAK-Sequence Error, or Implied NAK, during a specified period, exceeds the specified threshold, the QP will be handled as if the IB spec defined Retry Count was exceeded.
<b>Diagnostic Counters</b>	Added new diagnostic counters to evaluate the number of ICMC hits and misses for particular resources.

**Table 19 - Firmware Changes and New Feature History (Sheet 3 of 10)**

Feature/Change	Description
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”, on page 22</a>
<b>Rev. 14.22.1002</b>	
<b>Disable SL/diff Flow</b>	Added support for disable SL/diff flow to avoid performance degradation for single queue using multiple priorities. This functionality should not be used when DCB (PFC, ETS) is enabled.
<b>Software Reset Flow</b>	Software Reset Flow enables the device to recover from fatal errors. The flow includes software detection of a fatal error, automatic creations of an mstdump file for future debug by the software, and resetting of the device. The feature is enabled using an mlxconfig command.  <b>Note:</b> The flow is currently not supported on Multi host devices, Socket Direct devices and devices running management traffic (NCSI, MCTP).
<b>Steering Discard Packet Counters</b>	Any received packet which is dropped by the device is accounted for. To enable this functionality, the following counters were added to count the discard packets (per vport): <ul style="list-style-type: none"> <li>• <code>nic_receive_steering_discard</code>: Number of packets that completed the NIC Receive Flow Table steering, and were discarded because they did not match any flow in the final Flow Table.</li> <li>• <code>receive_discard_vport_down</code>: Number of packets that were steered to a VPort, and discarded because the VPort was not in a state to receive packets.</li> <li>• <code>transmit_discard_vport_down</code>: Number of packets that were transmitted by a vNIC, and discarded because the VPort was not in a state to transmit packets.</li> </ul>
<b>Virtual Functions (VF)</b>	Increased the number of VFs that can work with full VMQoS (8 TC) per PFs as follow: <ul style="list-style-type: none"> <li>• in dual port devices to 20 VFs</li> <li>• in single port devices to 58 VFs</li> </ul>
<b>Pause Frame Duration and XOFF Resend Time</b>	Increased the Pause Frame Duration and the XOFF Resend Time to the maximum value defined by the specification.
<b>PCI Relax Ordering</b>	mlxconfig configuration can now enable or disable forced PCI relaxed ordering in mkey_context. If this feature is enabled, the software per mkey configuration is ignored.
<b>vport Mirroring</b>	Packets are mirrored based on certain mirroring policy. The policy is set using the “set FTE command” that supports forward action in the ACL tables (ingress/egress). The firmware support the following destination list format: <ol style="list-style-type: none"> <li>1. new destination vport (analyzer)</li> <li>2. another Flow Table</li> </ol> <p>this way, the driver can forward the SX/RX packet related to the vport once it reaches the ACL table (forward it to the analyzer vport).</p>
<b>Resiliency: Special Error Event</b>	Firmware uses error events to monitor the health of core transport engines, both Rx and Tx, and to detect if a system hang occurred and was not cured by other error mechanisms. Upon such detection, events are sent to the driver to perform any required action (e.g., software reset).
<b>10GBaseT module</b>	Added support for 10GBaseT modules connected to a QSFP cage. <b>Note:</b> This connectivity supported was only tested with eNet's E10GSFPT-ENC 10GBase-T SFP+ device, and 10Gtek's ASF-10G-T when using firmware v14.22.1002
<b>QP's Creation Time</b>	Accelerated QP's creation time.

**Table 19 - Firmware Changes and New Feature History (Sheet 4 of 10)**

Feature/Change	Description
<b>SR-IOV LID based Routing Mode</b>	SR-IOV default routing mode is now LID based. The configuration change is available via mlxconfig tool. Note that in such mode, the VF will get its own LID, hence the GRH is not required.  <b>Note:</b> LID based routing support for vports is supported using SM v4.8.1
<b>Expansion ROM</b>	Added PXE and UEFI to additional ConnectX-4 Lxadapter cards. ConnectX-4 Lx now holds PXE, x86-UEFI and Arm-UEFI
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”, on page 22</a>
<b>Rev. 14.21.2010</b>	
<b>Query vPort Environments (Debug Counters)</b>	Debug counters are a group of counters that handle traffic performance issue related to firmware overhead in transport flow. The following are the additional counters added to this firmware version: <ul style="list-style-type: none"> <li>• current_q_under_processor_handle</li> <li>• total_q_under_processor_handle</li> <li>• qp_priority_update_flow</li> </ul>
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”, on page 22</a>
<b>Rev. 14.21.1000</b>	
<b>Receiver Signal Integrity Improvements</b>	Raised the network link only with phase greater than 15 ticks to improve signal integrity.
	Extended measurement test between 2 similar RX configurations.
	Moved the data path to use the second input buffer to improve signal integrity.
<b>Dynamic Bandwidth Allocation per Host</b>	Added support for new chassis MC OEM commands which allowed dynamically configuring rate limit and bandwidth share, per host and physical function.
<b>RoCE Dual Port Mode (a.k.a Multi-Port vHCA: MPV)</b>	Enables the usage of a dual port Virtual HCA (vHCA) to share RDMA resources (e.g., MR, CQ, SRQ, PDs) across the two Ethernet (RoCE) NIC network ports and display the NIC as a dual port device. For this feature to function properly, the following requirements must be met: <ul style="list-style-type: none"> <li>• Either the LAG or the Dual Port mode is enabled by the driver</li> <li>• Dual port device: both ports must be set as ETH</li> <li>• In ConnectX-4/ConnectX-4 Lx adapter cards, the maximum allowed number of VFs per PF is 32.</li> <li>• Function per port is enabled</li> </ul> <b>Note:</b> This feature is only supported in single host device
<b>DSCP</b>	Added QPDPM register to support dynamic mapping between DSCP and priority.
	Added trust level for QoS prioritization according to the DSCP or PCP.
	Added ingress buffer management for: <ul style="list-style-type: none"> <li>• ingress traffic mapping to a buffer according to priority</li> <li>• buffers sizes and lossless parameters</li> </ul>
<b>Steering Rules Rate Improvement</b>	Improved steering rules update rate to up to 50K rules per sec.
<b>Windows SR-IOV Enhanced eIPoIB</b>	Enabled Windows SR-IOV Enhanced eIPoIB (without Secure Connection) for Windows-over-Windows setups.
<b>Driver CR Dump</b>	crdump operation takes a snapshot of the device’s crspace dword-by-dword. It enables the driver to collect debug information upon firmware failure.

**Table 19 - Firmware Changes and New Feature History (Sheet 5 of 10)**

Feature/Change	Description
<b>Secured Firmware Update</b>	Secure Firmware Updates provides devices with the ability to verify digital signatures of new firmware binaries, in order to ensure that only officially approved versions are installed on the devices. <b>Note:</b> This feature is only available in adapter cards that support this feature.
<b>Cables</b>	Changed the default FEC mode for cables with attenuation 16 and below from RS to FC.
<b>ECN</b>	Enabled ECN (CongestionControl) by default for all priorities on Ethernet ports.
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”, on page 22</a>
<b>Rev. 14.20.1010</b>	
<b>DSCP</b>	Added trust level for QoS prioritization according to the DSCP or PCP. <b>[Beta]</b> Added ingress buffer management for: <ul style="list-style-type: none"> <li>ingress traffic mapping to a buffer according to priority</li> <li>buffers sizes and lossless parameters</li> </ul>
<b>Secured Firmware Updates</b>	<b>[Beta]</b> Secure Firmware Updates provides devices with the ability to verify digital signatures of new firmware binaries, in order to ensure that only officially approved versions are installed on the devices. <b>Note:</b> This feature is only available in adapter cards that support this feature.
<b>Relaxed Ordering</b>	<b>[Beta]</b> Added support for relaxed ordering write in memory keys.
<b>RDMA Counters</b>	Enhanced RDMA counter
<b>Outbound Buffer Overflow Mechanism</b>	Added the ability to enable and disable drop of packets in case of PCIe outbound buffer overflow. In addition added a counter for counting these drops.
<b>TLV for PCI class code</b>	Added 2 new per Host TLVs (see <a href="#">Table 21, “Supported Non-Volatile Configurations,” on page 37</a> )
<b>Fast Teardown</b>	Enables fast unloading driver by using Teardown HCA with <code>op_mode=1</code> (force_close). For further information, refer to the PRM.
<b>SFP Power Flow Improvement (level 2,1)</b>	Added support for SFP power class.
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”, on page 22</a>
<b>Rev. 14.18.2000</b>	
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”, on page 22</a>
<b>Rev. 14.18.1000</b>	
<b>RX Loss (BaseT link down indication)</b>	Added logical link indication in SFP to BaseT modules and disabled logical link when peer port is down.
<b>SFP Rate</b>	Added support for 10GbE in 25GbE SFP optical modules
<b>PDDR</b>	Enables mlxlink tool to collect data on the PHY link status and provides link down reasons and additional link related information.
<b>KR Tx Response</b>	Enabled TX configuration response and movement during Link Training in Ethernet.
<b>Phy Test mode</b>	Added support at lane rate of 12.89Gb.
<b>Head of Queue (HoQ) per TC</b>	Limits the amount of time a packet may head a Traffic Class (TC) transmission queue, without being transmitted. Stale packets are discarded. Active by default for TCs adhering to link level flow control

**Table 19 - Firmware Changes and New Feature History (Sheet 6 of 10)**

Feature/Change	Description
<b>User Access Region (UAR) 4KB Granularity Allocation</b>	UAR page size currently is set to 4KB and not according to what the system page size determines.
<b>No Driver NIC (NODNIC) Performance Improvement</b>	Improved performance of: <ul style="list-style-type: none"> <li>• Doorbell from User Access Region (UAR)</li> <li>• Clear interrupt from User Access Region (UAR)</li> </ul>
<b>Counters</b>	Added support for additional transport counters.
<b>On Demand Paging (ODP) DC</b>	Added ODP support for DC.
<b>Scatter to CQE on Sender for DC</b>	Enabled scatter-to-CQE for sent packets for DC.
<b>CQ modify</b>	Enabled moderation period modification in CQ modify command.
<b>VMQ: Rate limit per Function</b>	<b>[Beta]</b> Added support for minimum/maximum rate limit per vport in SR-IOV.
<b>Network traffic between UEFI-Shell and OS</b>	Enabled network traffic between UEFI-Shell and OS.
<b>non-RDMA capable VFs</b>	Enabled the PF to force disable RoCE for its VFs.
<b>PRM: Access Registers</b>	Added 2 new access registers: <ul style="list-style-type: none"> <li>• Management Capabilities Mask Register</li> <li>• Ports Capabilities Mask Register Fields</li> </ul> For further information, please refer to the PRM.
<b>Loopback Enabled/Disabled</b>	Enabled VNIC the control to enable/disable its local loopback traffic.
<b>RDMA RX Flow Table</b>	Added the option to open a receive RDMA Flow Table and to forward RoCE traffic to some destination QP.
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”, on page 22</a>
<b>Rev. 14.17.2020</b>	
<b>GENEVE &amp; IP-in-IP Stateless Offload</b>	<b>[Beta]</b> Added support for IP-in-IP and GENEVE network protocols encapsulated into IP frame (L2 tunneling). Encapsulation is suggested as a means to alter the normal IP routing for datagrams, by delivering them to an intermediate destination that would otherwise not be selected based on the (network part of the) IP Destination Address field in the original IP header. <b>Note:</b> For driver support, please see the Release Notes/User Manual of the relevant OS driver.
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”, on page 22</a>
<b>Rev. 14.17.1010</b>	
<b>Multi-Host LID Base Routing</b>	Added support for Multi-Host LID base routing. This feature requires a new OpenSM (v4.7.1 and above which comes with MLNX_OFED 3.3-2.0.0.0) with the following attributes: <ul style="list-style-type: none"> <li>• qos TRUE</li> <li>• lmc 2 (if there is no quad host in the fabric, you can set the lmc to 1)</li> <li>• virt_enabled 2</li> </ul> <b>Note:</b> Multi-Host LID base routing can be configured by the INI only. The default is 0
<b>Resilient RoCE</b>	Resilient RoCE is the ability to send RoCE traffic over a lossy network (a network without flow control enabled), without the need to enable flow control on the network. The ability is accomplished by enabling ECN on both the Switch and the Host.
<b>Multi-Host L3/L4 Classification</b>	Enables load balancing in the Multi PF Switch layer (MPFS) based on the L3/L4 headers

**Table 19 - Firmware Changes and New Feature History (Sheet 7 of 10)**

Feature/Change	Description
<b>Virtual Functions (VF) per Port</b>	<p>Increased the number of VFs from 64 to 95 per Physical Function (PF).</p> <p><b>Note:</b> When increasing the number of VFs, the following limitations must be taken into consideration:</p> <pre>server_total_bar_size &gt;= (num_pfs)*(2log_pf_uar_bar_size + 2log_vf_uar_bar_size*total_vfs) server_total_msix &gt;= (num_pfs)*(num_pf_msix + num_vfs_msix *total_vfs)</pre> <p><b>Note:</b> For the maximum number of VFs supported by your driver, please refer to your drivers' Release Notes or User Manual.</p>
<b>HCA Port Flap Counter</b>	Added support for Port Flap Counter.
<b>Fixed Buffer Size (KSM)</b>	Limits the buffer size for all entries to improve performance. KSM is used when associating Key Length My Virtual Address (KLMs) with fixed memory size.
<b>NULL Mkey</b>	This entry (null_mkey) is use to indicate non-present KLM/KSM entries. When accessing is, it causes the device to generate page fault event.
<b>Out-of-Band Online Firmware Update: Firmware Update over PLDM</b>	PLDM firmware burning is based on the DMTF spec DSP0267 (draft 9). The feature enables upgrading firmware and expansion ROM images using the PLDM protocol over MCTP (over PCIe). By doing so, a supporting BMC can query and upgrade the firmware without using OS based tools.
<b>New Group in Ports Performance Counters (PPCNT)</b>	Added a new physical layer statistics counters group. The new group includes BER counters, FEC error correction, clear time, and additional physical layer counters. For further information, please refer to the <a href="#">Ethernet Adapters Programming Manual (PRM)</a> .
<b>Permanent Link Up Mode</b>	<p>Enables the user to set a certain link up state for an unlimited period of time. This mode has 3 states:</p> <ul style="list-style-type: none"> <li>Aux power (standby)</li> <li>Reboot/boot/driver unloaded - the server is active and no driver is up</li> <li>Driver is up - at least one driver is up (the time between init HCA and teardown or FLR)</li> </ul>
<b>No Driver NIC (NODNIC) Performance Improvement</b>	Added support for Doorbell from User Access Region (UAR).
<b>SR-IOV: Rate Limit Per Function</b>	<b>[Beta]</b> Added support for maximum rate limit per function in SR-IOV.
<b>Firmware Resiliency: Suppress Pauses</b>	Allows the user to configure the adapter card to stop sending pauses after x when the receive port is unavailable (in a hang state).
<b>Performance Back-pressure Counters</b>	<b>[Beta]</b> Added support for new performance counters.
<b>Data Center Bridging Exchange (DCBX)</b>	DCBX is used by DCB devices to exchange configuration information with directly connected peers. DCBX uses Link Layer Discovery Protocol (LLDP) to exchange parameters between two link peers. For further information, please refer to the PRM.
<b>Access Register: Default Values Revert</b>	Allows network port registers to revert to their default values when the driver is restarted or the host is rebooted.
<b>Link up Modes</b>	<p>Added additional network link up modes. The new modes decide when to keep the network link up.</p> <p>The new modes are:</p> <ul style="list-style-type: none"> <li>keep_eth_link_up</li> <li>keep_ib_link_up</li> <li>keep_link_up_on_boot</li> <li>keep_link_up_on_standby</li> </ul>
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”, on page 22</a>

**Table 19 - Firmware Changes and New Feature History (Sheet 8 of 10)**

Feature/Change	Description
<b>Rev. 14.16.1020</b>	
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”</a> , on page 22
<b>Rev. 14.16.1006</b>	
<b>Explicit Congestion Notification (ECN)</b>	<b>[Beta]</b> Explicit Congestion Notification (ECN) is an extension to the Internet Protocol and to the Transmission Control Protocol. ECN allows end-to-end notification of network congestion without dropping packets.
<b>64 VFs per port</b>	Increased the number of VFs from 32 to 64 per PF. <b>Note:</b> When increasing the number of VFs, the following limitations must be taken into consideration: <code>server_total_bar_size &gt;= (num_pfs)*(2log_pf_uar_bar_size + 2log_vf_uar_bar_size*total_vfs)</code> <code>server_total_msix &gt;= (num_pfs)*(num_pf_msix + num_vfs_msix *total_vfs)</code>
<b>RoCE Link Aggregation (RoCE LAG)</b>	<b>[Beta]</b> RoCE Link Aggregation provides failover and link aggregation capabilities. In this mode, only one IB port, that represents the two physical ports, is exposed to the application layer. For further information, please refer to the PRM.
<b>OVS Offload</b>	Mellanox Accelerated Switching And Packet Processing (ASAP <sup>2</sup> ) Direct technology allows to offload OVS by handling OVS data-plane in Mellanox ConnectX-4 / ConnectX-4 Lx NIC hardware (Mellanox Embedded Switch or eSwitch) while maintaining OVS control-plane unmodified. The current actions supported by ASAP <sup>2</sup> Direct include packet parsing and matching, forward, drop along with VLAN push/pop or VXLAN encap/decap and HW based packet/byte flow statistics.
<b>Virtual Extensible LAN (VXLAN) encapsulation/decapsulation</b>	Virtual Extensible LAN (VXLAN) is a network virtualization technology that improves scalability problems associated with large cloud computing deployments. It tunnels Ethernet frames within Ethernet + IP + UDP frames. Mellanox implements VXLAN encapsulation and decapsulation in the hardware.
<b>Data Center Bridging Exchange (DCBX)</b>	<b>[Beta]</b> DCBX is used by DCB devices to exchange configuration information with directly connected peers. DCBX uses Link Layer Discovery Protocol (LLDP) to exchange parameters between two link peers. For further information, please refer to the PRM.
<b>FCS no scatter / FCS check</b>	Enables the user to control whether or not to scatter Frame Check Sequence (FCS) or to check FCS functionality.
<b>Packet Pacing</b>	<b>[Beta]</b> Send Queues (SQ/ Send queue of QP) may be individually rate limited, thus, allowing granular rate control over specific SW-defined flows. A rate-limited flow is allowed to transmit a few packets before its transmission rate is evaluated, and the next packet is scheduled for transmission accordingly.
<b>PRBS Patterns Generation and Tuning</b>	A new PHY test mode in which the device can generate different PRBS patterns for SerDes tuning purpose. For further information, please refer to PRM registers: PPAOS, PPTT, PPRT.
<b>Management Controller Transport Protocol (MCTP) over PCI</b>	Added support for MCTP host management over PCI
<b>OCBB / OCSD support after mlxfwreset</b>	Added support for OCBB/OCSD memory pointers restoration after mlxfwreset
<b>MCTP media migration</b>	Added support for MCTP media migration between SMBUS and PCI
<b>Cables</b>	Removed the RX amplitude configuration on some cable types

**Table 19 - Firmware Changes and New Feature History (Sheet 9 of 10)**

Feature/Change	Description
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”, on page 22</a>
<b>Rev. 14.14.2036</b>	
<b>Scatter FCS in RQ</b>	Enables software to scatter or strip FCS in RQ.
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”, on page 22</a>
<b>Rev. 14.14.1100</b>	
<b>CQE Time Stamping</b>	Keeps track of the creation of a packet. A time-stamping service supports assertions of proof that a datum existed before a particular time.
<b>Priority Flow Control (PFC)</b>	Applies pause functionality to specific classes of traffic on the Ethernet link.
<b>RDMA retransmission counters</b>	Custom port counters provide the user a clear indication about RDMA send/receive statistics and errors.
<b>Link Layer Discovery Protocol (LLDP)</b>	The Link Layer Discovery Protocol (LLDP) is a vendor-neutral Link Layer protocol in the Internet Protocol Suite used by network devices for advertising their identity, capabilities, and neighbors on a IEEE 802 LAN. The protocol is formally defined in IEEE 802.1AB.
<b>1GbE Link Speed</b>	ConnectX-4 Lx adapters now support 1Gb/s Ethernet connectivity in addition to 10GigE, 25GigE, 40GigE, 50GigE
<b>Flow Steering Counters</b>	Provides a clear indication of Flow Steering statistics and errors.
<b>WQE Inline Header</b>	The minimal amount of packet headers inlined in the WQE's Eth Segment.
<b>table-miss Flow</b>	A flow table may include a table-miss flow entry, which renders all Match Fields wildcards. If a packet does not match a flow entry in a flow table, this is a table miss. The behavior on a table miss depends on the table configuration. A table-miss flow entry in the flow table may specify how to process unmatched packets.
<b>Strided WQE User Space</b>	Striding RQ is a receive queue comprised by work queue elements (i.e. WQEs), where multiple packets of LRO segments (i.e. message) are written to the same WQE.
<b>SR-IOV (EN eSwitch &amp; RoCE)</b>	Single Root IO Virtualization (SR-IOV) is a technology that allows a physical PCIe device to present itself multiple times through the PCIe bus.
<b>Vector Calculation/ Erasure Coding Offload</b>	Uses the HCA for offloading erasure coding calculations.
<b>Firmware Image Time Stamping for Multi-Host Environment</b>	Enables the administrator to add a timestamp to the firmware they want to upgrade to avoid situations where one host tries to upgrade the firmware and another tries to downgrade; which may lead to two or more unnecessary server reboots. For further information, please refer to <a href="#">MFT User Manual</a> .
<b>Link params modification via access registers</b>	The change includes the following: 1. Changed port configuration which required link re-training (such as speed) 2. PAOS down 3. PAOS up This change, will cause the link to toggle and new configurations to take effect.
<b>Checksum Calculation on Image/Device</b>	Flint utility allows performing an MD5 checksum on the non-persistent sections of the firmware image. For further information, please refer to <a href="#">MFT User Manual</a> .
<b>Rev. 14.12.1240</b>	
<b>Bug Fixes</b>	See <a href="#">Section 4, “Bug Fixes History”, on page 22</a>
<b>Rev. 14.12.1100</b>	



**Table 19 - Firmware Changes and New Feature History (Sheet 10 of 10)**

Feature/Change	Description
<b>Port Link</b>	Reduced the port link-up time when negotiating according to Clause 73 (DME)
<b>Rev. 14.12.0780</b>	
<b>PCI</b>	<ul style="list-style-type: none"> <li>• PCIe Function Level Reset (FLR)</li> <li>• Power Management L2/L3 flow support</li> </ul>
<b>Ethernet Network</b>	<ul style="list-style-type: none"> <li>• Large Receive Offload (LRO)</li> <li>• Large Send Offload (LSO)</li> <li>• Receive Side Scaling (RSS)</li> <li>• Global Pause</li> <li>• RoCEv1.0/RoCEv2.0</li> <li>• Flow Steering</li> <li>• Sniffer Ethernet</li> <li>• Rate Limiter (at Beta level)</li> <li>• Multi packet WQE</li> <li>• Minimal Bandwidth Guarantee (ETS)</li> <li>• Explicit Congestion Notification (ECN)</li> <li>• Priority Flow Control (PFC)</li> </ul>
<b>PRM</b>	<ul style="list-style-type: none"> <li>• Strided SRQ</li> <li>• Self Loopback support</li> <li>• Transport Domain support</li> <li>• CQ2EQ remapping</li> <li>• Added support for the following commands:               <ul style="list-style-type: none"> <li>• MODIFY/QUERY_ESW_VPORT_CONTEXT</li> <li>• QUERY/MODIFY_CONG_STATUS</li> <li>• QUERY/MODIFY_CONG_PARAMS</li> <li>• QUERY_CONG_STATISTICS</li> <li>• ADD/DELETE_VXLAN_UDP_DPORT</li> </ul> </li> </ul>
<b>Virtualization</b>	<ul style="list-style-type: none"> <li>• VXLAN/NVGRE Stateless offload In this release, this feature is supported through Windows ONLY</li> <li>• SR-IOV EN (at Beta level)</li> </ul>
<b>Performance</b>	<ul style="list-style-type: none"> <li>• CQE zipping</li> </ul>
<b>Misc</b>	<ul style="list-style-type: none"> <li>• Wake-on-Lane/Standby</li> <li>• FlexBoot/UEFI support</li> </ul>
<b>Non-Volatile Configuration</b>	<ul style="list-style-type: none"> <li>• Non-Volatile Configuration (NVConfig). For the complete list, please refer to <a href="#">Section 8, on page 37</a>.</li> </ul>

## **6 PreBoot Drivers (FlexBoot/UEFI)**

### **6.1 FlexBoot Changes and New Features**

For further information, please refer to FlexBoot Release Notes

<https://docs.mellanox.com> --> Software --> Firmware Management --> PreBoot Drivers (FlexBoot/UEFI)

### **6.2 UEFI Changes and Major New Features**

For further information, please refer to UEFI Release Notes

<https://docs.mellanox.com> --> Software > Firmware Management --> PreBoot Drivers (FlexBoot/UEFI)

## 7 Unsupported Features and Commands

### 7.1 Unsupported Features

The following advanced feature are unsupported in the current firmware version:

**Table 20 - Unsupported Features**

Feature	ConnectX-4 Lx
The following service types: <ul style="list-style-type: none"> <li>• SyncUMR</li> <li>• Mellanox transport</li> <li>• RAW IPv6</li> </ul>	No
INT-A not supported for EQs only MSI-X	No
PCI VPD write flow (RO flow supported)	No
Streaming Receive Queue (STRQ) and collapsed CQ	No
Subnet Manager (SM) on VFs	No
DC in Multi-Host, SR-IOV, and Ethernet (RoCE)	No
RoCE LAG for VFs	No
RoCE LAG in Multi-Host/Socket-Direct	No
Mutlihost Ethernet	No

### 7.2 Unsupported Commands

- QUERY\_MAD\_DEMUX
- SET\_MAD\_DEMUX
- CREATE\_RQ - MEMORY\_RQ\_RMP
- MODIFY\_LAG\_ASYNC\_EVENT

## 8 Supported Non-Volatile Configurations

**Table 21 - Supported Non-Volatile Configurations**

Configuration	mlxconfig Parameter Name	Class	TLV ID
NV_MEMIC_CONF	MEMIC_BAR_SIZE	GLOBAL (0)	0x6
	MEMIC_SIZE_LIMIT		
NV_HOST_CHAINING_CONF	HOST_CHAINING_MODE		0x8
	HOST_CHAINING_DESCRIPTOR		
	HOST_CHAINING_TOTAL_BUFFER_SIZE		
NV_FLEX_PARS_CONF	FLEX_PARSER_PROFILE_ENABLE		0xe
	FLEX_IPV4_OVER_VXLAN_PORT		
NV_ROCE_1_5_CONF	ROCE_NEXT_PROTOCOL		0x10
NV_INTERNAL_RESOURCE_CONF	ESWITCH_HAIRPIN_DESCRIPTOR		0x13
	ESWITCH_HAIRPIN_TOT_BUFFER_SIZE		
NV_GLOBAL_PCI_CONF	NON_PREFETCHABLE_PF_BAR		0x80
	NUM_OF_VFS		
	SRIOV_EN		
	PF_LOG_BAR_SIZE		
	VF_LOG_BAR_SIZE		
	NUM_PF_MSIX		
	NUM_VF_MSIX		
NV_TPT_CONF	INT_LOG_MAX_PAYLOAD_SIZE		0x82
NV_POWER_CONF	SW_RECOVERY_ON_ERRORS		0x88
	RESET_WITH_HOST_ON_ERRORS		
	ADVANCED_POWER_SETTINGS		
NV_SW_OFFLOAD_CONFIG	CQE_COMPRESSION	0x10a	
	IP_OVER_VXLAN_EN		
	PCI_ATOMIC_MODE		
	LRO_LOG_TIMEOUT0		
	LRO_LOG_TIMEOUT1		
	LRO_LOG_TIMEOUT2		
	LRO_LOG_TIMEOUT3		
NV_IB_DC_CONF	LOG_DCR_HASH_TABLE_SIZE	0x190	
	DCR_LIFO_SIZE		
NV_VPI_LINK_TYPE	LINK_TYPE	PHYSICAL_PORT (2)	0x12

**Table 21 - Supported Non-Volatile Configurations**

Configuration	mlxconfig Parameter Name	Class	TLV ID
NV_ROCE_CC	ROCE_CC_PRIO_MASK	PHYSICAL_PORT (2)	0x107
	ROCE_CC_ALGORITHM		
NV_ROCE_CC_ECN	CLAMP_TGT_RATE_AFTER_TIME_INC	PHYSICAL_PORT (2)	0x108
	CLAMP_TGT_RATE		
	RPG_TIME_RESET		
	RPG_BYTE_RESET		
	RPG_THRESHOLD		
	RPG_MAX_RATE		
	RPG_AI_RATE		
	RPG_HAI_RATE		
	RPG_GD		
	RPG_MIN_DEC_FAC		
	RPG_MIN_RATE		
	RATE_TO_SET_ON_FIRST_CNP		
	DCE_TCP_G		
	DCE_TCP_RTT		
	RATE_REDUCE_MONITOR_PERIOD		
	INITIAL_ALPHA_VALUE		
	MIN_TIME_BETWEEN_CNPS		
	CNP_802P_PRIO		
CNP_DSCP			
NV_LLDP_NB_CONF	LLDP_NB_DCBX	PHYSICAL_PORT (2)	0x10a
	LLDP_NB_RX_MODE		
	LLDP_NB_TX_MODE		
NV_LLDP_NB_DCBX	DCBX_IEEE	PHYSICAL_PORT (2)	0x18e
	DCBX_CEE		
	DCBX_WILLING		
NV_KEEP_LINK_UP	KEEP_ETH_LINK_UP	PHYSICAL_PORT (2)	0x190
	KEEP_IB_LINK_UP		
	KEEP_LINK_UP_ON_BOOT		
	KEEP_LINK_UP_ON_STANDBY		
NV_QOS_CONF	NUM_OF_VL	PHYSICAL_PORT (2)	0x192
	NUM_OF_TC		
	NUM_OF_PFC		
NV_MPFS_CONF	DUP_MAC_ACTION	PHYSICAL_PORT (2)	0x196
	SRIOV_IB_ROUTING_MODE		
	IB_ROUTING_MODE		

**Table 21 - Supported Non-Volatile Configurations**

Configuration	mlxconfig Parameter Name	Class	TLV ID
NV_HCA_CONF	PCI_WR_ORDERING	HOST-FUNCTION (3)	0x112
	MULTI_PORT_VHCA_EN		
NV_EXTERNAL_PORT_CTRL	PORT_OWNER		0x192
	ALLOW_RD_COUNTERS		
	RENEG_ON_CHANGE		
	TRACER_ENABLE		
NV_ROM_BOOT_CONF2	IP_VER		0x195
	BOOT_UNDI_NETWORK_WAIT		
NV_ROM_UEFI_CONF	UEFI_HII_EN		0x196
NV_ROM_UEFI_DEBUG_LEVEL	BOOT_DBG_LOG		0x206
	UEFI_LOGS		
NV_ROM_BOOT_CONF1	BOOT_VLAN		0x221
	LEGACY_BOOT_PROTOCOL		
	BOOT_RETRY_CNT		
	BOOT_LACP_DIS		
	BOOT_VLAN_EN		
NV_ROM_IB_BOOT_CONF	BOOT_PKEY	0x222	
NV_PCI_CONF	ADVANCED_PCI_SETTINGS	HOST (7)	0x80
SAFE_MODE_CONF	SAFE_MODE_THRESHOLD		0x82
	SAFE_MODE_ENABLE		